We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

A Bio-Inspired Stereo Vision System for Guidance of Autonomous Aircraft

Richard J. D. Moore,
Saul Thurrowgood, Dean Soccol, Daniel Bland and Mandyam V. Srinivasan
*University of Queensland
Australia*

1. Introduction

Unmanned aerial vehicles (UAVs) are increasingly replacing manned systems in situations that are either too dangerous, too remote, or too difficult for manned aircraft to access. Modern UAVs are capable of accurately controlling their position and orientation in space using systems such as the Global Positioning System (GPS) and the Attitude and Heading Reference System (AHRS). However, they are unable to perform crucial guidance tasks such as obstacle avoidance, low-altitude terrain or gorge following, or landing in an uncontrolled environment using these systems only. For such tasks, the aircraft must be able to continuously monitor its surroundings. Active sensors, such as laser range finders or radar can be bulky, low-bandwidth, and stealth-compromising. Therefore, there is considerable benefit to be gained by designing guidance systems for UAVs that utilise passive sensing, such as vision. Over the last two decades, a significant amount of research has shown that biological visual systems can inspire novel, vision-based solutions to some of the challenges facing autonomous aircraft guidance. A recent trend in biologically inspired vision systems has been to exploit optical flow information for collision avoidance, terrain and gorge following, and landing. However, systems that rely on optical flow for extracting range information need to discount the components of optical flow that are induced by rotations of the aircraft. Furthermore, altitude cannot be controlled in a precise manner using measurements of optical flow only, as optical flow also depends upon the aircraft's velocity.

Stereo vision, on the other hand, allows the aircraft's altitude to be directly computed and controlled, irrespective of the attitude or ground speed of the aircraft, and independently of its rotations about the roll, pitch, and yaw axes. Additionally, we will show that a stereo vision system can also allow the computation and control of the orientation of the aircraft with respect to the ground. Stereo vision therefore provides an attractive approach to solving some of the problems of providing guidance for autonomous aircraft operating in low-altitude or cluttered environments.

In this chapter, we will explore how stereo vision may be applied to facilitate the guidance of an autonomous aircraft. In particular, we will discuss a wide-angle stereo vision system that has been tailored for the specific needs of aircraft guidance, such as terrain following, obstacle avoidance, and landing. Finally, results from closed-loop flight tests conducted using this system will be presented to demonstrate how stereo vision can be successfully utilised to provide guidance for an autonomous aircraft performing real-world tasks.

2. Relevant background

In this section we will briefly discuss the motivations for designing guidance systems for autonomous aircraft and also review some of the techniques used by state-of-the-art systems.

2.1 Unmanned aerial vehicles

Unmanned aerial vehicles (UAVs) have seen unprecedented levels of growth in both military and civilian application domains since their inception during World War I. So much so, in fact, that the Joint Strike Fighter¹, which is currently under production, is predicted to be the last manned aircraft produced by the US Armed Forces (Valavanis, 2007). The first pilotless aircraft were intended for use as aerial torpedoes. Today, however, autonomous or semi-autonomous fixed-wing aircraft, airships, or helicopters and vertical take-off and landing (VTOL) rotorcraft are increasingly being used for applications such as surveillance and reconnaissance, mapping and cartography, border patrol, inspection, military and defense missions, search and rescue, law enforcement, fire detection and fighting, agricultural and environmental imaging and monitoring, traffic monitoring, *ad hoc* communication networks, and extraterrestrial exploration, to name just a few.

The reason that UAVs are increasingly being preferred for these roles is that they are able to operate in situations that are either too dangerous, too remote, too dull, or too difficult for manned aircraft (Valavanis, 2007). Typically, today's UAVs are flown remotely by a human pilot. However, with the expanding set of roles there is an increasing need for UAVs to be able to fly with a degree of low-level autonomy, thus freeing up their human controllers to concentrate on high level decisions.

2.2 Short range navigation

Modern UAVs are capable of controlling their position and orientation in space accurately using systems such as the Global Positioning System (GPS) and Attitude and Heading Reference Systems (AHRS). This is sufficient when navigating over large distances at high altitude or in controlled airspaces. However, the expanding set of roles for UAVs increasingly calls for them to be able to operate in near-earth environments, and in environments containing 3D structures and obstacles. In such situations, the UAV must know its position in the environment accurately, which can be difficult to obtain using GPS due to occlusions and signal reflections from buildings and other objects. Additionally, the UAV must know *a priori* the 3D structure of the surrounding environment in order to avoid obstacles. Obviously such a scheme presents severe difficulties in situations where there is no foreknowledge of the 3D structure of the environment, or where this structure can change unpredictably. A more efficient approach would be for the aircraft to monitor its surroundings continuously during flight. The use of active proximity sensors such as ultrasonic or laser range finders, or radar has been considered for this purpose (Scherer et al., 2007). However, such systems can be bulky, expensive, stealth-compromising, high power, and low-bandwidth – limiting their utility for small-scale UAVs. Therefore, there is considerable benefit to be gained by designing guidance systems for UAVs that utilise passive sensing, such as vision.

2.3 Biological vision

The importance of vision for short range navigation was realised many decades ago. However, it is not until recently that vision-based guidance systems have been able to be

¹Lockheed Martin F-35 Lightning II.

demonstrated successfully onboard real robots outside controlled laboratory environments (see (DeSouza & Kak, 2002) for a review). The difficulty is that visual systems provide such a wealth of information about the surrounding environment and the self-motion of the vehicle, that it is a laborious task to extract the information necessary for robot guidance. For many animals, however, vision provides the primary sensory input for navigation, stabilisation of flight, detection of prey or predators, and interaction with other conspecifics. Insects in particular provide a good study model because they have seemingly developed efficient and effective visual strategies to overcome many of the challenges facing UAV guidance. For instance, the humble housefly is often able to outwit even the most determined swatter, despite its small brain and relatively simple nervous system. In fact, many flying insects have attained a level of skill, agility, autonomy, and circuit miniaturisation that greatly outperforms present day aerial robots (Franceschini, 2004).

Fig. 1. The optical flow, F , produced by an object as observed by an animal or robot in motion. Reproduced from (Hrabar et al., 2005), with permission.

Unlike vertebrates or humans, insects have immobile eyes with fixed-focus optics. Therefore, they cannot infer the distances to objects in the environment using cues such as the gaze convergence or refractive power required to bring an object into focus on the retina. Furthermore, compared with human eyes, the eyes of insects possess inferior spatial acuity, and are positioned much closer together with limited overlapping fields of view. Therefore, the precision with which insects could estimate range through stereopsis would be limited to relatively small distances (Srinivasan et al., 1993). Not surprisingly then, insects have evolved alternative strategies for overcoming the problems of visually guided flight. Many of these strategies rely on using image motion, or *optical flow*, generated by the insect's self-motion, to infer the distances to obstacles and to control various manoeuvres (Gibson, 1950; Nakayama & Loomis, 1974; Srinivasan et al., 1993). The relationship between optical flow and the range to objects in the environment is remarkably simple, and depends only upon the translational speed of the observer, the distance to the obstacle, and the azimuth of the obstacle with respect to the heading direction (Nakayama & Loomis, 1974) (see Fig. 1). The optical flow that is generated by the rotational motion of the insect does not encode any information on the range to objects and so must be discounted from the calculation. Alternatively, rotational movements of the vision system must be prevented and the optical flow measured when the vision system is undergoing pure translation.

For an observer translating at a speed v , and rotating at an angular velocity ω , the optical flow F , generated by a stationary object at a distance d , and angular bearing θ , is given by

$$F = \frac{v \times \sin(\theta)}{d} - \omega. \quad (1)$$

A significant amount of research over the past two decades has shown that biological vision systems can inspire novel, vision-based solutions to many of the challenges that must be overcome when designing guidance systems for autonomous aircraft (see (Srinivasan et al., 2004; Franceschini, 2004; Floreano et al., 2009) for reviews). It has been shown, for example, that honeybees use optical flow for negotiating narrow gaps and avoiding obstacles, regulating their flight speed and altitude, performing smooth landings, and estimating their distance flown (Srinivasan et al., 2000; Srinivasan & Zhang, 2004). A recent trend in biologically inspired vision systems for UAVs, therefore, has been to exploit optical flow information for collision avoidance, terrain and gorge following, and landing (Srinivasan et al., 2004; 2009; Barrows et al., 2003).

2.4 Existing bio-inspired vision-based guidance systems for UAVs

The magnitude of the optical flow gives a measure of the ratio of the aircraft's speed to its distance to objects in the environment. It has been demonstrated that both forward speed and altitude can be regulated using a single optical flow detector that was artificially maintained vertical (Ruffier & Franceschini, 2005). Altitude control for cruise flight has also been demonstrated onboard real UAV platforms (Barrows & Neely, 2000; Barrows et al., 2003; Green et al., 2003; 2004; Oh et al., 2004; Chahl et al., 2004) by regulating the ventral, longitudinal optical flow observed from the aircraft. While functional, the results of these early experiments were limited however, due to the failure to take the pitching motions of the aircraft into account and the passive or artificial stabilisation of roll. (Garratt & Chahl, 2008) also control altitude via optical flow and additionally correct for the pitching motions of the aircraft using an inertial measurement unit (IMU), but do not take the attitude of the aircraft into consideration. (Neumann & Bulthoff, 2001; 2002) use a similar strategy in simulation but regulate the UAV's attitude using colour gradients present in the simulated test environment. Using similar principles, (Thurrowgood et al., 2009; Todorovic & Nechyba, 2004) demonstrate methods for controlling UAV attitude based on the apparent orientation of the horizon and (Thakoor et al., 2003; 2002) use an attitude regulation scheme based on insect *ocelli*.

It has been proposed that insects, such as honeybees, navigate through narrow openings and avoid obstacles by balancing the optical flow observed on both sides of the body (Srinivasan et al., 1991), and by turning away from regions of high optical flow (Srinivasan & Lehrer, 1984; Srinivasan, 1993; Srinivasan & Zhang, 1997). Similar strategies have been employed by (Zufferey & Floreano, 2006; Zufferey et al., 2006; Green et al., 2004; Green, 2007; Oh et al., 2004; Hrabar & Sukhatme, 2009) to demonstrate lateral obstacle avoidance in aircraft. (Beyeler et al., 2007; Beyeler, 2009; Zufferey et al., 2008) steer to avoid obstacles in three dimensions and additionally incorporate rate gyroscopes and an anemometer to account for the motions of the aircraft and the measure the airspeed of the aircraft respectively. The study of insect behaviour has also revealed novel strategies which may be used to control complex flight manoeuvres. It has been observed that as honeybees land they tend to regulate their forward speed proportionally to their height such that the optical flow produced by the landing surface remains constant (Srinivasan et al., 2000). As their height approaches zero, so does their forward speed, ensuring a safe, low speed at touch down for the bee. Similar strategies have

been employed by (Beyeler, 2009; Chahl et al., 2004; Green et al., 2003; 2004; Oh et al., 2004) to demonstrate autonomous take-off and landing of small scale UAVs.

Obviously, therefore, measuring the optical flow produced by the motion of the aircraft through the environment is a viable means of providing guidance information for an autonomous aircraft, as is attested by the success of the approaches described above. However, extracting the necessary information from the observed optical flow is not without its difficulties.

2.5 Stereo vision

Optical flow is inherently noisy, and obtaining dense and accurate optical flow images is computationally expensive. Additionally, systems that rely on optical flow for extracting range information need to discount the components of optical flow that are induced by rotations of the aircraft, and use only those components that are generated by the translational motion of the vehicle (see Equation 1). This either requires an often noisy, numerical estimate of the roll, pitch, and yaw rates of the aircraft, or additional apparatus for their explicit measurement, such as a three-axis gyroscope. Furthermore, the range perceived from a downward facing camera or optical flow sensor is not only dependent upon altitude and velocity, but also the aircraft's attitude. This is particularly relevant to fixed-wing aircraft in which relatively high roll and pitch angles are required to perform rapid manoeuvres. A method for overcoming these shortcomings is described in (Beyeler et al., 2006), however the technique proposed there is too limited to be implemented in practice as it fails to include the roll angle of the aircraft. Finally, as with all optical flow-based approaches, to retrieve an accurate estimate of range, the ground-speed of the aircraft must be decoupled from the optical flow measurement. In practice this requires additional sensors, such as high-precision GPS or a Pitot tube. Moreover, in the case of the latter, the variable measured is actually airspeed, which would lead to incorrect range estimates in all but the case of low altitude flight in still air.

Stereo vision, on the other hand, allows the aircraft's altitude to be directly measured and controlled, irrespective of the attitude or ground-speed of the aircraft, and independently of its rotations about the roll, pitch, and yaw axes. Additionally, for stereo systems the visual search is constrained to a single dimension, hence reducing the complexity and increasing the accuracy of the computation. Furthermore, we will show that a wide-angle stereo vision system can also allow the computation and control of the orientation of the aircraft with respect to the ground. Stereo vision therefore provides an attractive approach to solving some of the problems of providing guidance for autonomous aircraft operating in low-altitude or cluttered environments.

Stereo vision systems have previously been designed for aircraft. An altitude regulation scheme is presented by (Roberts et al., 2002; 2003) who use a downwards facing stereo system to measure the height of an aircraft, although they require that the attitude is regulated via an onboard IMU. (Hrabar & Sukhatme, 2009) also require that the attitude of their aircraft is externally regulated and they utilise a combined stereo and optical flow approach to navigate urban canyons and also avoid frontal obstacles. Wide-angle stereo vision systems have also been investigated (Thurrowgood et al., 2007; Tisse et al., 2007), but they have rarely been tailored to the specific needs of aircraft guidance, such as terrain and gorge following, obstacle avoidance, and landing. In this chapter, we describe a stereo vision system that is specifically designed to serve these requirements.

3. A bio-inspired stereo vision system for UAV guidance

In this section we introduce a wide-angle stereo vision system that is tailored to the specific needs of aircraft guidance. The concept of the vision system is inspired by biological vision systems and its design is intended to reduce the complexity of extracting appropriate guidance commands from the visual data. The vision system was originally designed to simplify the computation of optical flow, but this property also makes it well suited to functioning as a coaxial stereo system. The resultant vision system therefore makes use of the advantages of stereo vision whilst retaining the simplified control schemes enabled by the bio-inspired design of the original vision system. In this section we will discuss the design, development, and implementation of the vision system, and also present results that demonstrate how stereo vision may be utilised to provide guidance for an autonomous aircraft.

3.1 Conceptual design

The concept of the vision system is best described by considering an assembly in which a camera views a specially shaped reflective surface (the mirror). As well as increasing the field of view (FoV) of the camera, the profile of the mirror is designed such that equally spaced points on the ground, on a line parallel to the camera's optical axis, are imaged to points that are equally spaced in the camera's image plane. This has the effect of removing the perspective distortion (and therefore the distortion in image motion) that a camera experiences when viewing a horizontal plane that stretches out to infinity in front of the aircraft. The mapping produced by the mirror is illustrated in Fig. 2. It is clear that equal distances along the ground, parallel to the optical axis of the system, are mapped to equal distances on the image plane, validating the design of the mirror. The full derivation of the mirror profile is given in (Srinivasan et al., 2006).

Fig. 2. Illustration of the imaging properties of the mirror. The raw image as viewed by the camera (left), and the remapped image (right). The dark line indicates viewing directions at 90° to the camera's optical axis. Reproduced from (Srinivasan et al., 2006).

The special geometric remapping afforded by the mirror means that, for a given vehicle speed, the motion in the camera's imaging plane of the image of an object in the environment is inversely proportional to the radial distance of that object from the optical axis of the vision system. Therefore, surfaces of constant image motion, when reprojected into the environment, are cylindrical, as is illustrated in Fig. 3. This property makes the system particularly useful for aircraft guidance. For any given aircraft speed, the maximum image velocity that is observed

Fig. 3. Illustration of the clear-space mapping provided by the vision system. Reproduced from (Srinivasan et al., 2006).

in the remapped image specifies the radius of a cylinder of space in front of the aircraft, through which collision-free flight can occur. This approach of characterising the collision-free space in front of the aircraft by a virtual cylinder simplifies the problem of determining in advance whether an intended flight trajectory through the environment will be collision-free, and of making any necessary course corrections to facilitate this.

Fig. 4. Schematic illustration of the conceptual stereo vision system, surface of constant disparity, and collision-free cylinder. Reproduced from (Moore et al., 2010).

Now consider a system in which two such camera-mirror assemblies are arranged coaxially, as illustrated in Fig. 4. Each camera views the environment through a mirror that has the imaging properties described above. It follows that the pixel disparity, D_{pixel} , produced by a point imaged in both cameras is inversely proportional to the radial distance, d_{radial} , of that point from the common optical axis of the two camera-mirror assemblies. The relationship is given by

$$D_{pixel} = \frac{d_{baseline} \times h_{image}}{r} \times \frac{1}{d_{radial}}, \tag{2}$$

where $d_{baseline}$ is the stereo baseline, h_{image} is the vertical resolution of the remapped images and r , the forward viewing factor, is the ratio of the total forward viewing distance to the height of the aircraft.

The first term in Equation 2 is simply a constant which depends on the system configuration (see Table 1 for representative values). Therefore, the maximum image disparity, in a

Stereo baseline ($d_{baseline}$)	200mm
Remap image cols / rows (h_{img})	200px / 288px
Vertical FoV	0° to 74° from vertical
Horizontal FoV	−100° to 100° from vertical
Forward viewing factor (r)	3.5
Detectable disparity (D_{pixel})	0px to 10px
Operational altitude (d_{radial})	1.5m ~ 50m+

Table 1. System parameters and their typical values.

given stereo pair, directly defines the radius of the collision-free cylinder that surrounds the optical axis, independent of the speed of the aircraft. Thus, a simple control loop may be implemented in which the aircraft is repelled from objects penetrating the notional flight cylinder required by the aircraft for collision-free flight. Furthermore, the image disparity will be one dimensional only, thereby reducing the complexity of the computation. The system is therefore well suited to providing real-time information for visual guidance in the context of tasks such as terrain and gorge following, obstacle detection and avoidance, and landing.

3.2 Hardware and implementation

In recent implementations of the vision system (Moore et al., 2009; 2010), the function of the specially shaped mirrors is simulated using software lookup tables. This requires calibrated camera-lens assemblies in order to generate the lookup tables but reduces the physical bulk and cost of the system and avoids aberrations due to imperfections in the mirror surfaces. The software remapping process is illustrated in Fig. 5. In this example an image of a rendered scene is captured through a rectilinear lens with a 120° FoV. The shaded area of the raw image is unwrapped and transformed to produce the remapped image. A comparison with Fig. 2 indicates that the image remapped in software shares the same properties with the image remapped in hardware (simulation), as expected.

Fig. 5. Illustration of the software remapping process. The shaded area in the raw image (a) is remapped to (b). Reproduced from (Moore et al., 2009).

The centre of the raw image (Fig. 5) is not remapped because, in this region, equal distances along the ground plane project onto infinitesimally small distances on the cameras' image

planes. Therefore, if remapped, the resolution in this region would be negligible. Using the representative system parameters listed in Table 1, the central unmapped region is a conic section surrounding the optical axis with radius approximately 16° . Thus, by situating the cameras coaxially, the field of view of the system is not compromised as this region is not remapped in any case². The outer diameter of the area to be remapped is limited by the FoV of the camera-lens assemblies.

Fig. 6. (a) Implementation of the stereo vision system from (Moore et al., 2009) and (b) mounted on the aircraft. Reproduced from (Moore et al., 2009).

The two cameras are rigidly mounted in a coaxial stereo configuration (Fig. 6) to minimise measurement errors resulting from relative motion between the two camera-lens assemblies during flight. In the implementation described in (Moore et al., 2009), we use high resolution video cameras (PGR GRAS-20S4M) equipped with wide-angle fish-eye lenses (Fujinon FE185C057HA-1), to provide good spatial resolution whilst maintaining a large FoV. However, in (Moore et al., 2010) we found it necessary to use lightweight miniature fish-eye lenses (Sunex DSL215) to reduce the vibration-induced motion of the lenses relative to the camera sensors, without compromising the FoV. In both implementations the stereo cameras are synchronised to within $125\mu\text{s}$ across the IEEE 1394b bus interface.

Each camera-lens assembly is calibrated according to the generic camera model described in (Kannala & Brandt, 2006). The two assemblies are also calibrated as a stereo pair so that any rotational misalignment can be compensated for during the remapping process. This stereo calibration is performed on the calibrated raw images from each camera. The sum of absolute pixel differences between the two stereo images of a distant scene is minimised by applying a three-degrees-of-freedom (3DOF) rotation to one of the camera models. The optimal corrective 3DOF rotation is found using the NLOpt library (Johnson, 2009) implementation of the BOBYQA algorithm (Powell, 2009).

The image disparity between stereo pairs is computed using an algorithm based on the sum of absolute pixel differences (SAD) between images and is implemented using the Intel Integrated Performance Primitives library (Intel, 2009). To remove low frequency image intensity gradients, which can confuse the SAD algorithm, the remapped images are convolved with a Scharr filter kernel before the disparity is computed. The SAD algorithm gives the image offset for a window surrounding each pixel for which the computed SAD score is a minimum. An equiangular fit, as described in (Shimizu & Okutomi, 2003), is applied to the minimum and neighbouring SAD scores for each window to obtain sub-pixel disparity estimates. Incorrect matches are rejected by re-computing the disparity for the

²This is also true if physical mirrors are used, as this region would be obscured by the self-reflection of each camera. This phenomenon is not visible in Fig. 2, as the camera body is not rendered in this case.

reverse image order and discarding disparities that differ from the initial estimate. This bidirectional technique is effective at rejecting mismatches but doubles the execution time of the algorithm. Further details on the specific implementations can be found in (Moore et al., 2009; 2010).

Stereo disparities are extracted from the remapped images via the process described above. The radial distances to objects in the environment can then be calculated from the stereo disparities according to Equation 2. Values for the system parameters can be selected depending on the desired operational range envelope, the required resolution of the range estimates, or the computational time available for processing disparities. Representative values used in (?) for flight testing are listed in Table 1.

Fig. 7. The aircraft used for flight testing was a Super Frontier Senior-46 (wingspan 2040mm), modified so that the engine and propeller assembly is mounted above the wing.

All image processing and higher order functions are performed on the aircraft (Fig. 7) by the onboard computer (Digital-Logic MSM945 which incorporates an Intel Core2 Duo 1.5GHz processor). Flight commands are sent continuously to the aircraft's control surfaces through an interface which allows a ground-based human pilot to select between computer controlled autonomous flight and radio controlled manual flight.

3.3 Range testing

The performance and accuracy of the stereo system were evaluated in (Moore et al., 2009) using an artificially textured arena. A cropped image of the arena as viewed by the front camera is displayed in Fig. 8. The texture used to line the arena is composed of black circles of varying diameter (65mm $\rightarrow$ 150mm) on a white background. The dimensions of the arena are 3200mm $\times$ 2350mm $\times$ 1150mm. The stereo rig was positioned in the centre of the arena with the optical axis parallel to the longest dimension of the arena. Also displayed in Fig. 8 is the remapped view of the testing arena overlayed with the computed stereo disparities. It can be seen that the disparity vectors have a constant magnitude in any given column³. This verifies the expected result – that image disparity between stereo pairs depends only on the radial distance of the viewed points from the optical axis.

Estimates of the radial distance to observed points can be calculated from the disparities via Equation 2. In order to quantify the accuracy of the stereo system, the relationship between the estimated radial distance to the arena and the viewing angle was plotted against the actual relationship (Fig. 9). The actual relationship was calculated from the known geometry of the arena. The viewing angle in this case corresponds to the vertical elevation of the viewing ray –

³Note that the disparity vectors that correspond to the rear wall of the arena are large in magnitude compared with surrounding vectors as these areas lie close to the optical axis of the vision system. These vectors are omitted from the following discussion to simplify analysis.

Fig. 8. (a) Cropped image of the testing arena as seen by the front camera and (b) the same view of the arena after remapping. The computed stereo disparities are overlaid in white. The disparity vectors have been scaled to aid visualisation. Reproduced from (Moore et al., 2009).

Fig. 9. Profile of the estimated radial distance to the arena wall and floor (blue) shown alongside the actual radial distance at each viewing angle (black). Error bars represent $\pm 2\sigma$ at each viewing angle. Reproduced from (Moore et al., 2009).

i.e. points that lie in the same column in the remapped image (Fig. 8) share the same viewing elevation. The error in the estimated radial distance at each viewing angle in Fig. 9 thus represents the variance from the multiple estimates at each viewing elevation. It can be seen that the errors in the estimated radial distances are most significant for viewing elevations that correspond to where the walls of the arena join the floor. This is a result of the non-zero size of the window used to compute the stereo disparity. A window size larger than one pixel would be expected to cause an underestimation of the radial distance to the corners of the arena, where surrounding pixels correspond to closer surfaces. Indeed this is observed in Fig. 9. Similarly, it would be expected to observe a slight overestimation in the radial distance to the arena floor directly beneath the vision system, where surrounding pixels correspond to surfaces that are further away, and this is also observed in Fig. 9.

The data presented in Fig. 9 is computed from a single typical stereo pair and is unfiltered, however a small number of points were rejected during the disparity computation. Small errors in the reprojected viewing angles may arise from inaccurate calibration of the camera-lens assemblies but are presumed to be negligible in this analysis. Therefore, the total error in the reconstruction can be specified as the error in the radial distance to the arena

at each viewing angle. The standard deviation of this error, measured from approximately 2.5×10^4 reprojected points, was $\sigma = 3.5 \times 10^{-2}\text{m}$, with very little systematic bias (systematic variance amongst points at the same viewing elevation). Represented as a percentage of the estimated radial distance at each viewing angle, the absolute (unsigned) reprojection error was calculated as having a mean of 1.2% and a maximum of 5.6%. This error is a direct consequence of errors in the computed stereo disparities.

4. UAV attitude and altitude stabilisation

In section 3, a closed-loop control scheme using a stereo vision system was described in which the aircraft was repelled from objects that penetrate a notional flight cylinder surrounding the flight trajectory. This control scheme provides an effective collision avoidance strategy for an autonomous UAV and also provides the ability to demonstrate behaviours such as terrain and gorge following. In this section we will show that the attitude and altitude of the aircraft with respect to the ground may also be measured accurately using the same stereo vision system. This enhancement provides for more precise attitude and altitude control during terrain following, and also allows for other manoeuvres such as constant altitude turns and landing. We will present results from recent closed-loop flight tests that demonstrate the ability of this vision system to provide accurate and real-time control of the attitude and altitude of an autonomous aircraft.

4.1 Estimating attitude and altitude

If it is assumed that the ground directly beneath and in front of the aircraft can be modelled as a planar surface, then the attitude of the aircraft can be measured with respect to the plane normal. Also, the altitude of the aircraft can be specified as the distance from the nodal point of the vision system to the ground plane, taken parallel to the plane normal. The attitude and altitude of the aircraft can therefore be measured from the parameters of a planar model fitted to the observed disparity points.

Two approaches for fitting the model ground plane to the observed disparities have been considered in this study. In (Moore et al., 2009) we fit the model ground plane in disparity space and in (Moore et al., 2010) we apply the fit in 3D space. The first approach is more direct but perhaps unintuitive. Given the known optics of the vision system, the calibration parameters of the cameras and the attitude and altitude of the aircraft carrying the vision system, the magnitudes and directions of the view rays that emanate from the nodal point of the vision system and intersect with the ideal infinite ground plane can be calculated. By reformulating the ray distances as radial distances from the optical axis of the vision system, the ideal disparities may be calculated via Equation 2. Thus, the disparity surface that should be measured by the stereo vision system at some attitude and altitude above an infinite ground plane can be predicted. Conversely, given the measured disparity surface, the roll, pitch, and height of the aircraft with respect to the ideal ground plane can be estimated by iteratively fitting the modelled disparity surface to the measurements. This is a robust method for estimating the attitude and altitude of the aircraft because the disparity data is used directly, hence the data points and average error will be distributed evenly over the fitted surface.

In order to fit the modelled disparity surface to the observed data, we must parameterise the disparity model using the roll, pitch, and height of the aircraft above the ground plane. We start by calculating the intersection between our view vectors and the ideal plane. A point on a line can be parameterised as $\mathbf{p} = t\hat{\mathbf{v}}$, where in our case $\hat{\mathbf{v}}$ is a unit view vector and t is

the distance to the intersection point from the origin (nodal point of the vision system), and a plane can be defined as $\mathbf{p} \cdot \hat{\mathbf{n}} + d = 0$. Solving for t gives

$$t = \frac{-d|\mathbf{v}|}{\mathbf{v} \cdot \hat{\mathbf{n}}}. \quad (3)$$

Now, in the inertial frame⁴, our ideal plane will remain stationary (our aircraft will rotate), so we define $\hat{\mathbf{n}} = [0 \ 0 \ -1]$. Therefore, $d = d_{\text{height}}$ is the distance from the ideal plane to the origin and, conversely, the height of the aircraft above the ground plane. So, making the substitutions,

$$t = \frac{-d_{\text{height}}|\mathbf{v}|}{\mathbf{v} \cdot [0 \ 0 \ -1]} = \frac{d_{\text{height}}|\mathbf{v}|}{\mathbf{v}_z}, \quad (4)$$

thus we must only find the z component of our view vector in the inertial frame.

In the camera frame, the z axis is parallel with the optical axis and the x and y axes are parallel with the rows and columns of the raw images respectively. Thus, our view vector is defined by the viewing angle, ν , taken around the positive z axis from the positive x axis, and the forward viewing ratio, r . Thus, $\mathbf{v}_{\text{cam}} = [\cos \nu \ \sin \nu \ r]$. To find the view vector in the inertial frame, we first transpose our view vector from the camera frame to the body frame, $\mathbf{v}_{\text{body}} = [r \ -\cos \nu \ -\sin \nu]$ (our cameras are mounted upside down), and then we rotate from the body frame to the inertial frame (we neglect yaw since our ideal ground plane does not define a heading direction).

$$\mathbf{R}_x(\phi) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \phi & -\sin \phi \\ 0 & \sin \phi & \cos \phi \end{bmatrix}, \text{ represents a rolling motion, } \phi, \text{ about the } x \text{ axis.}$$

$$\mathbf{R}_y(\theta) = \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}, \text{ represents a pitching motion, } \theta, \text{ about the } y \text{ axis.}$$

Therefore,

$$\mathbf{R}_{\text{body} \rightarrow \text{world}}(\phi, \theta) = \mathbf{R}_y(\theta) \times \mathbf{R}_x(\phi) = \begin{bmatrix} \cos \theta & \sin \theta \sin \phi & \sin \theta \cos \phi \\ 0 & \cos \phi & -\sin \phi \\ -\sin \theta & \cos \theta \sin \phi & \cos \theta \cos \phi \end{bmatrix},$$

and

$$\mathbf{v}_{\text{world}} = \mathbf{R}_{\text{body} \rightarrow \text{world}}(\phi, \theta) \times \mathbf{v}_{\text{body}}. \quad (5)$$

Now, we are only interested in $\mathbf{v}_z$, the z component of the view vector, $\mathbf{v}_{\text{world}}$. Therefore, multiplying out Equation 5 gives

$$\mathbf{v}_z^i = -\cos(\theta) \sin(\nu^i + \phi) - r^i \sin(\theta), \quad (6)$$

where we have included the superscript i to indicate that this is the view vector corresponding to the i^{th} pixel in the remapped image. Substituting Equation 6 back into Equation 4 gives

$$t^i = \frac{d_{\text{height}}|\mathbf{v}^i|}{-\cos(\theta) \sin(\nu^i + \phi) - r^i \sin(\theta)}, \quad (7)$$

⁴We use the common aeronautical North-East-Down (NED) inertial frame.

where t^i is the direct ray distance to the ideal ground plane along a particular view vector. Now, the stereo vision system actually measures the radial distance to objects from the optical axis. Therefore to convert t in Equation 7 from ray distance to radial distance, we drop the scale factor $|\mathbf{v}|$. So finally, substituting Equation 7 back into Equation 2, we get the expected disparity surface measured by the stereo vision system for a particular attitude and altitude above an ideal ground plane,

$$D_{pixel}^i = \frac{d_{baseline} \times h_{image}}{r_{tot}} \times \frac{1}{d_{height}} \times \left[-\cos(\theta) \sin(\nu^i + \phi) - r^i \sin(\theta) \right], \quad (8)$$

where the first term is a system constant as described before, and the radial distance has been replaced by d_{height} , the vertical height (in the inertial frame) of the aircraft above the ideal ground plane. The bracketed term describes the topology of the disparity surface and depends on the roll, ϕ , and pitch, θ , of the aircraft as well as two parameters ν^i and r^i , that determine the viewing angles in the x and z (camera frame) planes respectively for the i^{th} pixel in the remapped image.

Fig. 10. Attitude and altitude of the aircraft during an outdoor test as estimated via fitting the disparity surface (black, solid), and fitting the 3D point cloud (blue, dashed). Also shown for comparison (red, dotted) are the roll and pitch angles as reported by an IMU and the depth measurement reported by an acoustic sounder. Frames were captured at approximately 12Hz.

In order to obtain the roll, pitch, and height of the aircraft, we minimise the sum of errors between Equation 8 and the measured disparity points using a non-linear, derivative-free optimisation algorithm. Currently, we use the NLOpt library (Johnson, 2009) implementation of the BOBYQA algorithm (Powell, 2009). This implementation typically gives minimisation times in the order of 10ms (using $\sim 6 \times 10^3$ disparities on a 1.5GHz processor). To analyse the performance of this approach, an outdoor test was conducted in which the lighting and texture

conditions were not controlled. The attitude and altitude estimates computed using this approach are shown in Fig. 10 plotted alongside the measurements from an IMU (MicroStrain 3DM-GX2) and a depth sounder, which were installed onboard the aircraft to provide distinct measurements of the attitude and altitude. It can be seen that the visually estimated motions of the aircraft correlate well with the values used for comparison.

The second approach for determining the attitude and altitude of the aircraft with respect to an ideal ground plane is to re-project the disparity points into 3D coordinates relative to the nodal point of the vision system and fit the ideal ground plane in 3D space. While this procedure does not sample data points uniformly in the plane, it leads to a single-step, non-iterative optimisation that offers the advantage of low computational overheads and reliable real-time operation. This is the approach taken in (Moore et al., 2010) to achieve real-time, closed-loop flight.

To re-project the disparity points into 3D space, we use the radial distances computed directly from the disparities via Equation 2,

$$\mathbf{p}^i = \frac{d_{rad}^i}{\sin \alpha^i} \cdot \hat{\mathbf{u}}^i, \quad (9)$$

where $\mathbf{p}^i$ is the reprojected location of the i^{th} pixel in 3D coordinates relative to the nodal point of the vision system, $\hat{\mathbf{u}}^i$ is the unit view vector for the i^{th} pixel (derived from the calibration parameters of the cameras) and α^i is the angle between the view vector and the optical axis.

Fig. 11. 3D reconstruction of the test arena. Measurements are in metres relative to the nodal point of the vision system. Reproduced from (Moore et al., 2009).

The radial distances computed from the stereo image pair of the test arena (seen in Fig. 8) were used to reconstruct the arena in 3D space (Fig. 11). It was found in (Moore et al., 2009) that the mean error in the radial distance estimates was approximately 1.2% for the test conducted in the indoor arena. It can be seen from Fig. 11 that this error leads to an accurate 3D reconstruction of the simple test environment. However, this reprojection error is directly attributable to the errors in the computed stereo disparities – which are approximately constant for any measurable disparity. Therefore, for the system parameters used during range testing (see (Moore et al., 2009)), the mean radial distance error of 1.2% actually indicates a mean error in the computed stereo disparities of approximately $\frac{1}{4}$ pixel.

The (approximately constant) pixel noise present in the disparity measurements means that at higher altitudes the range estimates will be increasingly noisy. This phenomena is responsible

for the maximum operational altitude listed in Table 1, for at altitudes higher than this maximum, the disparity generated by the ground is less than the mean pixel noise. Thus, for altitudes within the operational range, fitting the ideal ground plane model to the reprojected 3D point cloud, rather than fitting the model to the disparities directly, results in less well constrained estimates of the orientation of the ideal plane, and hence less well constrained estimates of the attitude and altitude of the aircraft. However, it can be seen from Fig. 10 that this approach is still a viable means of estimating the state of the aircraft, particularly at altitudes well below the operational limit of the system. Furthermore, this approach results in an optimisation that is approximately two orders of magnitude faster than the first approach discussed above. This is because the optimisation can be performed in a single-step using a least-squares plane fit on the 3D point cloud. In (Moore et al., 2010) we use a least-squares algorithm from the WildMagic library (Geometric Tools, 2010) and achieve typical optimisation times in the order of $< 1\text{ms}$ (using $\sim 6 \times 10^3$ reprojected points on a 1.5GHz processor).

Applying the planar fit in 3D space therefore offers lower computational overheads at the cost of reduced accuracy in the state estimates. However, the least-squares optimisation may be implemented within a RANSAC⁵ framework to reject outliers and improve the accuracy of the state estimation. This is the approach taken in (Moore et al., 2010) to achieve closed-loop control of an aircraft performing time-critical tasks such as low-altitude terrain following.

4.2 Closed-loop terrain following

During flight, the stereo vision system discussed in this chapter can provide real-time estimates of the attitude and altitude of an aircraft with respect to the ground plane using the methods described above. However, for autonomous flight, the aircraft must also generate control commands appropriate for the desired behaviour. In (Moore et al., 2010), we use cascaded proportional-integral-derivative (PID) feedback control loops to generate the flight commands whilst attempting to minimise the error between the visually estimated altitude and attitude and their respective setpoints. The closed-loop control scheme is depicted in Fig. 12. Roll and pitch are controlled independently and so full autonomous control is achieved using two feedback control subsystems. Additionally, within each control subsystem, multiple control layers are cascaded to improve the stability of the system.

Fig. 12. Block diagram illustrating the closed-loop control scheme used for closed-loop flight testing. Reproduced from (Moore et al., 2010).

⁵RANdom Sample Consensus. An iterative method for estimating function parameters in the presence of outliers.

Fig. 13. Visually estimated height (black, solid) and pitch angle (blue, dashed) during a segment of flight. Also shown is a scaled binary trace (red, shaded) that indicates the periods of autonomous control, during which the aircraft was programmed to hold an altitude of 10m AGL. Reproduced from (Moore et al., 2010).

The control subsystem for stabilising the roll of the aircraft comprises two cascaded PID controllers. The highest level controller measures the error in the roll angle of the aircraft and delivers an appropriate roll rate command to the lower level controller, which implements the desired roll rate. The pitch control subsystem functions identically to the roll subsystem, although it includes an additional cascaded PID controller to incorporate altitude stabilisation. Shown in Fig. 12, aircraft altitude is regulated by the highest level PID controller, which feeds the remainder of the pitch control subsystem. Measurements of the absolute attitude and altitude of the aircraft are made by the stereo vision system and are used to drive all other elements of the closed-loop control system. Low level control feedback for the roll rate and pitch rate is provided by an onboard IMU. The multiple control layers allow the aircraft to be driven towards a particular altitude, pitch angle, and pitch rate simultaneously. This allows for stable control without the need for accurately calibrated integral and derivative gains. It is observed that a more responsive control system may be produced by collapsing the absolute angle and rate controllers into a single PID controller for each subsystem (where the rate measurements from the IMU are used by the derivative control component). However, the closed-loop data presented in this section was collected using the control system described by Fig. 12.

The closed-loop performance of the vision system was evaluated in (Moore et al., 2010) by piloting the test aircraft (Fig. 7) in a rough racetrack pattern. During each circuit the aircraft was piloted to attain an abnormal altitude and attitude, and then automatic control was engaged for a period of approximately 5s – 10s. A quantitative measure of the performance of the system was then obtained by analysing the ability of the aircraft to restore the set attitude and altitude of 0° roll angle and 10m above ground level (AGL) respectively. This procedure was repeated 18 times during a test flight lasting approximately eight minutes. A typical segment of flight (corresponding to 380s ~ 415s in Fig. 15) during which the aircraft made two autonomous passes is shown in Figs. 13 & 14. It can be seen that on both passes, once autonomous control was engaged, the aircraft was able to attain and hold the desired attitude and altitude within approximately two seconds. It can also be seen that the visually estimated roll angle closely correlates with the measurement from the IMU throughout the flight segment. Temporary deviations between the estimated roll and pitch angles and the

Fig. 14. Visually estimated roll angle (black, solid) during a segment of flight. For comparison, the roll angle reported by an onboard IMU is shown (blue, dashed). Also shown is a scaled binary trace (red, shaded) that indicates the periods of autonomous control, during which the aircraft was programmed to hold a roll angle of 0° with respect to the ground plane. Reproduced from (Moore et al., 2010).

values reported by the IMU are to be expected, however, due to the inherent difference between the measurements performed by the stereo vision system, which measures attitude with respect to the local orientation of the ground plane, and the IMU, which measures attitude with respect to gravity.

The visually estimated altitude of the aircraft throughout the full flight test is displayed in Fig. 15. It can be seen that in every autonomous pass the aircraft was able to reduce the absolute error between its initial altitude and the setpoint (10m AGL), despite initial altitudes varying between 5m and 25m AGL. The performance of the system was measured by considering two metrics: the time that elapsed between the start of each autonomous segment and the aircraft first passing within one metre of the altitude setpoint; and the average altitude of the aircraft during the remainder of each autonomous segment (*i.e.* not including the initial response phase). These metrics were used to obtain a measure of the response time and steady-state accuracy of the system respectively. From the data presented in Fig. 15, the average response time of the system was calculated as $1.45s \pm 1.3s$, where the error bounds represent 2σ from the 18 closed-loop trials. The relatively high variance of the average response time is due to the large range of initial altitudes. Using the second metric defined above, the average unsigned altitude error was calculated as $6.4 \times 10^{-1}m$ from approximately 92s of continuous segments of autonomous terrain following. These performance metrics both indicate that the closed-loop system is able to quickly respond to sharp adjustments in altitude and also that the system is able to accurately hold a set altitude, validating its use for tasks such as autonomous terrain following.

5. Conclusions

This chapter has introduced and described a novel, wide-angle stereo vision system for the autonomous guidance of aircraft. The concept of the vision system is inspired by biological vision systems and its design is intended to reduce the complexity of extracting appropriate guidance commands from visual data. The vision system takes advantage of the accuracy and reduced computational complexity of stereo vision, whilst retaining the simplified control

Fig. 15. The visually estimated altitude (black, solid) of the aircraft during the flight test. Also shown is a scaled binary trace (red, dashed) that indicates the periods of autonomous control, during which the aircraft was programmed to hold an altitude of 10m AGL. Reproduced from (Moore et al., 2010).

schemes enabled by its bio-inspired design. Two coaxially aligned video cameras are used in conjunction with two wide-angle lenses to capture stereo imagery of the environment, and a special geometric remapping is employed to simplify the computation of range. The maximum disparity, as measured by this system, defines a collision-free cylinder surrounding the optical axis through which the aircraft can fly unobstructed. This system is therefore well suited to providing visual guidance for an autonomous aircraft in the context of tasks such as terrain and gorge following, obstacle detection and avoidance, and take-off and landing. Additionally, it was shown that this stereo vision system is capable of accurately measuring and representing the three dimensional structure of simple environments, and two control schemes were presented that facilitate the measurement of the attitude and altitude of the aircraft with respect to the local ground plane. It was shown that this information can be used by a closed-loop control system to successfully provide real-time guidance for an aircraft performing autonomous terrain following. The ability of the vision system to react quickly and effectively to oncoming terrain has been demonstrated in closed-loop flight tests. Thus, the vision system discussed in this chapter demonstrates how stereo vision can be effectively and successfully utilised to provide visual guidance for an autonomous aircraft.

6. Acknowledgments

This work was supported partly by US Army Research Office MURI ARMY-W911NF041076, Technical Monitor Dr Tom Doligalski, US ONR Award N00014-04-1-0334, ARC Centre of Excellence Grant CE0561903, and a Queensland Smart State Premier's Fellowship. The authors are associated with the Queensland Brain Institute & School of Information Technology and Electrical Engineering, University of Queensland, St Lucia, Australia and ARC Centre of Excellence in Vision Science, Australia. Finally, sincere thanks to Mr. David Brennan, who owns and maintains the airstrip at which the flight testing was done.

7. References

- Barrows, G. L., Chahl, J. S. & Srinivasan, M. V. (2003). Biologically inspired visual sensing and flight control, *The Aeronautical Journal* 107(1069): 159–168.
- Barrows, G. L. & Neely, C. (2000). Mixed-mode VLSI optic flow sensors for in-flight control of a micro air vehicle, *Proc. SPIE*, Vol. 4109, pp. 52–63.
- Beyeler, A. (2009). *Vision-based control of near-obstacle flight*, PhD thesis, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland.
- Beyeler, A., Mattiussi, C., Zufferey, J.-C. & Floreano, D. (2006). Vision-based altitude and pitch estimation for ultra-light indoor aircraft, *Proc. IEEE International Conference on Robotics and Automation (ICRA'06)*, pp. 2836–2841.
- Beyeler, A., Zufferey, J.-C. & Floreano, D. (2007). 3D vision-based navigation for indoor microflyers, *Proc. IEEE International Conference on Robotics and Automation (ICRA'07)*, Roma, Italy.
- Chahl, J. S., Srinivasan, M. V. & Zhang, S. W. (2004). Landing strategies in honeybees and applications to uninhabited airborne vehicles, *The International Journal of Robotics Research* 23(2): 101–110.
- DeSouza, G. N. & Kak, A. C. (2002). Vision for mobile robot navigation: A survey, 24(2).
- Floreano, D., Zufferey, J.-C., Srinivasan, M. V. & Ellington, C. P. (2009). *Flying Insects and Robots*, Springer. In press.
- Franceschini, N. (2004). Visual guidance based on optic flow: A biorobotic approach, *Journal of Physiology* 98: 281–292.
- Garratt, M. A. & Chahl, J. S. (2008). Vision-based terrain following for an unmanned rotorcraft, *Journal of Field Robotics* 25: 284–301.
- Geometric Tools (2010). Wildmagic library.
URL:<http://www.geometrictools.com/LibMathematics/Approximation/Approximation.html>
- Gibson, J. J. (1950). *The Perception of the Visual World*, Houghton Mifflin.
- Green, W. E. (2007). *A Multimodal Micro Air Vehicle for Autonomous Flight in Near-Earth Environments*, PhD thesis, Drexel University, Philadelphia, PA.
- Green, W. E., Oh, P. Y. & Barrows, G. L. (2004). Flying insect inspired vision for autonomous aerial robot maneuvers in near-earth environments, *Proc. IEEE International Conference on Robotics and Automation (ICRA'04)*, New Orleans, LA.
- Green, W. E., Oh, P. Y., Sevcik, K. & Barrows, G. (2003). Autonomous landing for indoor flying robots using optic flow, *Proc. ASME International Mechanical Engineering Congress*, Washington, D.C.
- Hrabar, S. & Sukhatme, G. (2009). Vision-based navigation through urban canyons, *Journal of Field Robotics* 26(5): 431–452.
- Hrabar, S., Sukhatme, G. S., Corke, P., Usher, K. & Roberts, J. (2005). Combined optic-flow and stereo-based navigation of urban canyons for a UAV, *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'05)*, Edmonton, Canada.
- Intel (2009). Integrated performance primitives library.
URL:<http://software.intel.com/sites/products/collateral/hpc/ipp/ippindepth.pdf>
- Johnson, S. G. (2009). The NLOpt nonlinear-optimization package.
URL:<http://ab-initio.mit.edu/nlopt>
- Kannala, J. & Brandt, S. S. (2006). A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses, 28(8): 1335–1340.
- Moore, R. J. D., Thurrowgood, S., Bland, D., Soccol, D. & Srinivasan, M. V. (2009). A stereo vision system for UAV guidance, *Proc. IEEE International Conference on Intelligent*

- Robots and Systems (IROS'09)*, St Louis, MO.
- Moore, R. J. D., Thurrowgood, S., Bland, D., Soccol, D. & Srinivasan, M. V. (2010). UAV altitude and attitude stabilisation using a coaxial stereo vision system, *Proc. IEEE International Conference on Robotics and Automation (ICRA'10)*, Anchorage, AK.
- Nakayama, K. & Loomis, J. M. (1974). Optical velocity patterns, velocity-sensitive neurons, and space perception: A hypothesis, *Perception* 3(1): 63–80.
- Neumann, T. & Bulthoff, H. H. (2001). Insect inspired visual control of translatory flight, *Proc. 6th European Conference on Artificial Life (ECAL'01)*, Prague, Czech Republic.
- Neumann, T. & Bulthoff, H. H. (2002). Behaviour oriented vision for biomimetic flight control, *Proc. EPSRC/BBSRC International Workshop on Biologically Inspired Robotics*, Bristol, UK.
- Oh, P. Y., Green, W. E. & Barrows, G. L. (2004). Neural nets and optic flow for autonomous micro-air-vehicle navigation, *Proc. ASME International Mechanical Engineering Congress and Exposition*, Anaheim, CA.
- Powell, M. (2009). The BOBYQA algorithm for bound constrained optimization without derivatives, *Cambridge NA Report NA2009/06*, University of Cambridge, Cambridge, Reino Unido.
- Roberts, J. M., Corke, P. I. & Buskey, G. (2002). Low-cost flight control system for a small autonomous helicopter, *Proc. Australasian Conference on Robotics and Automation (ACRA'02)*, Auckland, New Zealand.
- Roberts, J. M., Corke, P. I. & Buskey, G. (2003). Low-cost flight control system for a small autonomous helicopter, *Proc. IEEE International Conference on Robotics and Automation (ICRA'03)*, Taipei, Taiwan.
- Ruffier, F. & Franceschini, N. (2005). Optic flow regulation: the key to aircraft automatic guidance, *Robotics and Autonomous Systems* 50: 177–194.
- Scherer, S., Singh, S., Chamberlain, L. & Saripalli, S. (2007). Flying fast and low among obstacles, *Proc. IEEE International Conference on Robotics and Automation (ICRA'07)*, Roma, Italy.
- Shimizu, M. & Okutomi, M. (2003). Significance and attributes of subpixel estimation on area-based matching, *Systems and Computers in Japan* 34(12).
- Srinivasan, M. V. (1993). How insects infer range from visual motion, *Reviews of Occulomotor Research* 5: 139–156.
- Srinivasan, M. V. & Lehrer, M. (1984). Temporal acuity of honeybee vision: behavioural studies using moving stimuli, *Journal of Comparative Physiology* 155: 297–312.
- Srinivasan, M. V., Lehrer, M., Kirchner, W. H. & Zhang, S. W. (1991). Range perception through apparent image speed in freely-flying honeybees, *Visual Neuroscience* 6: 519–535.
- Srinivasan, M. V., Thurrowgood, S. & Soccol, D. (2006). An optical system for guidance of terrain following in UAV's, *Proc. IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS'06)*, Sydney, Australia, pp. 51–56.
- Srinivasan, M. V., Thurrowgood, S. & Soccol, D. (2009). From flying insects to autonomously navigating robots, *16(3)*: 59–71.
- Srinivasan, M. V. & Zhang, S. (2004). Visual motor computations in insects, *Annual Review of Neuroscience* 27: 679–696.
- Srinivasan, M. V. & Zhang, S. W. (1997). Visual control of honeybee flight, *Orientation and Communication in Arthropods* 84: 95–113.
- Srinivasan, M. V., Zhang, S. W., Chahl, J. S., Barth, E. & Venkatesh, S. (2000). How honeybees make grazing landings on flat surfaces, *Biological Cybernetics* 83(3): 171–183.

- Srinivasan, M. V., Zhang, S. W., Chahl, J. S., Stange, G. & Garratt, M. (2004). An overview of insect inspired guidance for application in ground and airborne platforms, *Proc. Inst. Mech. Engrs. Part G* 218: 375–388.
- Srinivasan, M. V., Zhang, S. W. & Chandrashekara, K. (1993). Evidence for two distinct movement-detecting mechanisms in insect vision, *Naturwissenschaften* 80: 38–41.
- Thakoor, S., Chahl, J., Srinivasan, M. V., Young, L., Werblin, F., Hine, B. & Zornetzer, S. (2002). Bioinspired engineering of exploration systems for NASA and DoD, *Artificial life* 8(4): 357–369.
- Thakoor, S., Zornetzer, S., Hine, B., Chahl, J. & Stange, G. (2003). Bioinspired engineering of exploration systems: a horizon sensor/attitude reference system based on the dragonfly ocelli for mars exploration applications, *Journal of Robotic Systems* 20(1): 35–42.
- Thurrowgood, S., Soccol, D., Moore, R. J. D., Bland, D. & Srinivasan, M. V. (2009). A vision based system for attitude estimation of UAVs, *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'09)*, St Louis, MO.
- Thurrowgood, S., Stuerzl, W., Soccol, D. & Srinivasan, M. V. (2007). A panoramic stereo imaging system for aircraft guidance, *Proc. Ninth Australasian Conference on Robotics and Automation (ACRA'07)*, Brisbane, Australia.
- Tisse, C. L., Frank, O. & Durrant-Whyte, H. (2007). Hemispherical depth perception for slow-flyers using coaxially aligned fisheye cameras, *Proc. International Symposium on Flying Insects and Robots*, Ascona, Switzerland, p. 123.
- Todorovic, S. & Nechyba, M. C. (2004). A vision system for intelligent mission profiles of micro air vehicles, *IEEE Transactions on Vehicular Technology* 53(6): 1713–1725.
- Valavanis, K. P. (2007). *Advances in Unmanned Aerial Vehicles: State of the Art and the Road to Autonomy*, Springer.
- Zufferey, J.-C., Beyeler, A. & Floreano, D. (2008). Near-obstacle flight with small UAVs, *Proc. International Symposium on Unmanned Aerial Vehicles (UAV'08)*, Orlando, FL.
- Zufferey, J.-C. & Floreano, D. (2006). Fly-inspired visual steering of an ultralight indoor aircraft, 22: 137–146.
- Zufferey, J.-C., Klapotocz, A., Beyeler, A., Nicoud, J.-D. & Floreano, D. (2006). A 10-gram microflyer for vision-based indoor navigation, *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'06)*, Beijing, China.

Advances in Theory and Applications of Stereo Vision

Edited by Dr Asim Bhatti

ISBN 978-953-307-516-7

Hard cover, 352 pages

Publisher InTech

Published online 08, January, 2011

Published in print edition January, 2011

The book presents a wide range of innovative research ideas and current trends in stereo vision. The topics covered in this book encapsulate research trends from fundamental theoretical aspects of robust stereo correspondence estimation to the establishment of novel and robust algorithms as well as applications in a wide range of disciplines. Particularly interesting theoretical trends presented in this book involve the exploitation of the evolutionary approach, wavelets and multiwavelet theories, Markov random fields and fuzzy sets in addressing the correspondence estimation problem. Novel algorithms utilizing inspiration from biological systems (such as the silicon retina imager and fish eye) and nature (through the exploitation of the refractive index of liquids) make this book an interesting compilation of current research ideas.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Richard J. D. Moore (2011). A Bio-Inspired Stereo Vision System for Guidance of Autonomous Aircraft, *Advances in Theory and Applications of Stereo Vision*, Dr Asim Bhatti (Ed.), ISBN: 978-953-307-516-7, InTech, Available from: <http://www.intechopen.com/books/advances-in-theory-and-applications-of-stereo-vision/a-bio-inspired-stereo-vision-system-for-guidance-of-autonomous-aircraft>

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2011 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.

IntechOpen

IntechOpen